Backport193 - Backport #8080

Segfault in rb_fd_set

03/13/2013 07:48 AM - jonleighton (Jon Leighton)

Status:	Closed	
Priority:	Normal	
Assignee:	usa (Usaku NAKAMURA)	

Description

I have experienced a segfault with Ruby 2 during an IO.select call:

See https://travis-ci.org/jonleighton/spring/jobs/5393025 or https://gist.github.com/jonleighton/5147785 to see the crash output.

I cannot reproduce on a different version of Linux (Fedora). However I was able to reproduce by downloading a VM image of the Travis CI environment and running the code on there (see http://pivotallabs.com/debugging-travis-builds/ for how to do that).

I tried to produce a simple script to reproduce, but without success. I also tried to build Ruby 2 with debugging symbols, but this did not produce the crash. I'm not sure why - perhaps related to compiler optimisations.

I found a workaround for the crash with https://github.com/jonleighton/spring/commit/c8a7afdd3238ef88bffc2c8f56baa21042400e15.

Related issues:

Has duplicate Ruby master - Bug #6653: 1.9.2/1.9.3 exhibit SEGV with many thr...	**Closed**	**06/27/2012**

Associated revisions

Revision 92b367e0 - 03/16/2013 05:07 AM - kosaki (Motohiro KOSAKI)

- thread.c: disabled _FORTIFY_SOURCE for avoid to hit glibc bug. [Bug #8080] [ruby-core:53349]
- test/ruby/test_io.rb (TestIO#test_io_select_with_many_files): test for the above.

git-svn-id: svn+ssh://ci.ruby-lang.org/ruby/trunk@39775 b2dd03c8-39d4-4d8f-98ff-823fe69b080e

Revision 39775 - 03/16/2013 05:07 AM - kosaki (Motohiro KOSAKI)

- thread.c: disabled _FORTIFY_SOURCE for avoid to hit glibc bug. [Bug #8080] [ruby-core:53349]
- test/ruby/test_io.rb (TestIO#test_io_select_with_many_files): test for the above.

Revision 39775 - 03/16/2013 05:07 AM - kosaki (Motohiro KOSAKI)

- thread.c: disabled _FORTIFY_SOURCE for avoid to hit glibc bug. [Bug #8080] [ruby-core:53349]
- test/ruby/test_io.rb (TestIO#test_io_select_with_many_files): test for the above.

Revision 39775 - 03/16/2013 05:07 AM - kosaki (Motohiro KOSAKI)

- thread.c: disabled _FORTIFY_SOURCE for avoid to hit glibc bug. [Bug #8080] [ruby-core:53349]
- test/ruby/test_io.rb (TestIO#test_io_select_with_many_files): test for the above.

Revision 39775 - 03/16/2013 05:07 AM - kosaki (Motohiro KOSAKI)

- thread.c: disabled _FORTIFY_SOURCE for avoid to hit glibc bug. [Bug #8080] [ruby-core:53349]
- test/ruby/test_io.rb (TestIO#test_io_select_with_many_files): test for the above.

Revision 39775 - 03/16/2013 05:07 AM - kosaki (Motohiro KOSAKI)

- thread.c: disabled _FORTIFY_SOURCE for avoid to hit glibc bug. [Bug #8080] [ruby-core:53349]
- test/ruby/test_io.rb (TestIO#test_io_select_with_many_files): test for the above.

Revision 39775 - 03/16/2013 05:07 AM - kosaki (Motohiro KOSAKI)

- thread.c: disabled _FORTIFY_SOURCE for avoid to hit glibc bug. [Bug #8080] [ruby-core:53349]
- test/ruby/test_io.rb (TestIO#test_io_select_with_many_files): test for the above.

Revision ccb9fb0b - 03/20/2013 01:34 PM - nagachika (Tomoyuki Chikanaga)

merge revision(s) 39772,39773: [Backport #8080]

```
    * configure.in: check struct timeval exist or not.

    * include/ruby/missing.h (struct timeval): check HAVE_STRUCT_TIMEVAL
      properly. and don't include sys/time.h if struct timeval exist.

    * file.c: include sys/time.h explicitly.

    * random.c: ditto.

    * thread_pthread.c: ditto.

    * time.c: ditto.

    * ext/date/date_strftime.c: ditto.

    * include/ruby/missing.h (struct timespec): include <sys/time.h>
```

git-svn-id: svn+ssh://ci.ruby-lang.org/ruby/branches/ruby_2_0_0@39838 b2dd03c8-39d4-4d8f-98ff-823fe69b080e

Revision 552817e5 - 03/20/2013 01:36 PM - nagachika (Tomoyuki Chikanaga)

merge revision(s) 39774: [Backport #8080]

```
    * include/ruby/missing.h (__syscall): moved to...

    * io.c: here. because __syscall() is only used from io.c.

    * include/ruby/missing.h: move "#include <sys/type.h>" to ....

    * include/ruby/intern.h: here. because it was introduced for
      fixing NFDBITS issue. [ruby-core:05179].
```

git-svn-id: svn+ssh://ci.ruby-lang.org/ruby/branches/ruby_2_0_0@39839 b2dd03c8-39d4-4d8f-98ff-823fe69b080e

Revision 7ff076d5 - 03/20/2013 01:37 PM - nagachika (Tomoyuki Chikanaga)

merge revision(s) 39775: [Backport #8080]

```
    * thread.c: disabled _FORTIFY_SOURCE for avoid to hit glibc bug.
      [Bug #8080] [ruby-core:53349]

    * test/ruby/test_io.rb (TestIO#test_io_select_with_many_files):
      test for the above.
```

git-svn-id: svn+ssh://ci.ruby-lang.org/ruby/branches/ruby_2_0_0@39840 b2dd03c8-39d4-4d8f-98ff-823fe69b080e

Revision d4e03967 - 03/20/2013 01:54 PM - nagachika (Tomoyuki Chikanaga)

merge revision(s) 39160: [Backport #8080]

```
    * configure.in: don't define ARCH_FLAG="-march=i486" if it causes
      compilation problem.
```

git-svn-id: svn+ssh://ci.ruby-lang.org/ruby/branches/ruby_2_0_0@39843 b2dd03c8-39d4-4d8f-98ff-823fe69b080e

Revision 78a2c1a9 - 03/20/2013 01:56 PM - nagachika (Tomoyuki Chikanaga)

merge revision(s) 39162,39163: [Backport #8080]

```
    * configure.in: change CFLAGS temporally to test
      ARCH_FLAG="-march=i486".
```

git-svn-id: svn+ssh://ci.ruby-lang.org/ruby/branches/ruby_2_0_0@39844 b2dd03c8-39d4-4d8f-98ff-823fe69b080e

Revision 7d772fbd - 03/20/2013 01:57 PM - nagachika (Tomoyuki Chikanaga)

merge revision(s) 39198: [Backport #8080]

```
    * configure.in: move the test for -march=i486 just after
      RUBY_UNIVERSAL_ARCH/RUBY_DEFAULT_ARCH.
```

git-svn-id: svn+ssh://ci.ruby-lang.org/ruby/branches/ruby_2_0_0@39845 b2dd03c8-39d4-4d8f-98ff-823fe69b080e

Revision c8755ba6 - 03/20/2013 02:01 PM - nagachika (Tomoyuki Chikanaga)

merge revision(s) 39174: [Backport #8080]

```
    * configure.in: move header files check to the beginning of
      "header and library section".
      test rlim_t with sys/types.h and sys/time.h for MirOS BSD.
      sys/types.h and sys/time.h is gurded by #ifdef and the above move
      is required for this change.
```

git-svn-id: svn+ssh://ci.ruby-lang.org/ruby/branches/ruby_2_0_0@39846 b2dd03c8-39d4-4d8f-98ff-823fe69b080e

Revision 9b0861fb - 03/20/2013 02:03 PM - nagachika (Tomoyuki Chikanaga)

merge revision(s) 39200: [Backport #8080]

```
    * configure.in: move OS specific header/function knowledge before
      automatic header tests.
```

git-svn-id: svn+ssh://ci.ruby-lang.org/ruby/branches/ruby_2_0_0@39847 b2dd03c8-39d4-4d8f-98ff-823fe69b080e

Revision 68a2dd54 - 03/20/2013 02:06 PM - nagachika (Tomoyuki Chikanaga)

merge revision(s) 39777: [Backport #8080]

```
    mingw build fix

    * configure.in: struct timeval is defined in winsock2.h on mingw.

    * include/ruby/missing.h: include time.h for time_t, and sys/time.h
```

for timeval and timespec.

git-svn-id: svn+ssh://ci.ruby-lang.org/ruby/branches/ruby_2_0_0@39848 b2dd03c8-39d4-4d8f-98ff-823fe69b080e

Revision 4346f5d6 - 03/20/2013 02:07 PM - nagachika (Tomoyuki Chikanaga)

merge revision(s) 39779: [Backport #8080]

```
    Makefile.sub: fix mswin build

    * win32/Makefile.sub (config.h): fix mswin build, also VC has time.h
```

and struct timeval.

git-svn-id: svn+ssh://ci.ruby-lang.org/ruby/branches/ruby_2_0_0@39849 b2dd03c8-39d4-4d8f-98ff-823fe69b080e

Revision 7631f28c - 03/20/2013 02:10 PM - nagachika (Tomoyuki Chikanaga)

merge revision(s) 39781,39783: [Backport #8080]

```
    missing.h: build fix

    * include/ruby/missing.h: include time.h and sys/time.h iff needed,
```

but excepct for sys/time.h on linux to get rid of glibc bug.
* include/ruby/missing.h: removed **linux**. it's unnecessary.

git-svn-id: svn+ssh://ci.ruby-lang.org/ruby/branches/ruby_2_0_0@39850 b2dd03c8-39d4-4d8f-98ff-823fe69b080e

Revision b97e9255 - 03/28/2013 10:10 AM - usa (Usaku NAKAMURA)

merge revision(s) 39772,39773,39774,39775,39777,39779,39781,39783: [Backport #8080]

```
    * configure.in: check struct timeval exist or not.

    * include/ruby/missing.h (struct timeval): check HAVE_STRUCT_TIMEVAL
      properly. and don't include sys/time.h if struct timeval exist.

    * file.c: include sys/time.h explicitly.

    * random.c: ditto.

    * thread_pthread.c: ditto.

    * time.c: ditto.
```

```
    * ext/date/date_strftime.c: ditto.

    * include/ruby/missing.h (struct timespec): include <sys/time.h>

    * include/ruby/missing.h (__syscall): moved to...

    * io.c: here. because __syscall() is only used from io.c.

    * include/ruby/missing.h: move "#include <sys/type.h>" to ....

    * include/ruby/intern.h: here. because it was introduced for
      fixing NFDBITS issue. [ruby-core:05179].

    * thread.c: disabled _FORTIFY_SOURCE for avoid to hit glibc bug.
      [Bug #8080] [ruby-core:53349]

    * test/ruby/test_io.rb (TestIO#test_io_select_with_many_files):
      test for the above.

    * include/ruby/missing.h: removed __linux__. it's unnecessary.
```

git-svn-id: svn+ssh://ci.ruby-lang.org/ruby/branches/ruby_1_9_3@39985 b2dd03c8-39d4-4d8f-98ff-823fe69b080e

Revision 39985 - 03/28/2013 10:10 AM - usa (Usaku NAKAMURA)

merge revision(s) 39772,39773,39774,39775,39777,39779,39781,39783: [Backport #8080]

```
* configure.in: check struct timeval exist or not.

* include/ruby/missing.h (struct timeval): check HAVE_STRUCT_TIMEVAL
  properly. and don't include sys/time.h if struct timeval exist.

* file.c: include sys/time.h explicitly.

* random.c: ditto.

* thread_pthread.c: ditto.

* time.c: ditto.

* ext/date/date_strftime.c: ditto.

* include/ruby/missing.h (struct timespec): include <sys/time.h>

* include/ruby/missing.h (__syscall): moved to...

* io.c: here. because __syscall() is only used from io.c.

* include/ruby/missing.h: move "#include <sys/type.h>" to ....

* include/ruby/intern.h: here. because it was introduced for
  fixing NFDBITS issue. [ruby-core:05179].

* thread.c: disabled _FORTIFY_SOURCE for avoid to hit glibc bug.
  [Bug #8080] [ruby-core:53349]

* test/ruby/test_io.rb (TestIO#test_io_select_with_many_files):
  test for the above.

* include/ruby/missing.h: removed __linux__. it's unnecessary.
```

Revision 9c215874 - 03/29/2013 04:23 AM - usa (Usaku NAKAMURA)

- include/ruby/missing.h: fixed merge mistake of r39985. [Backport #8080]

git-svn-id: svn+ssh://ci.ruby-lang.org/ruby/branches/ruby_1_9_3@39995 b2dd03c8-39d4-4d8f-98ff-823fe69b080e

Revision 39995 - 03/29/2013 04:23 AM - usa (Usaku NAKAMURA)

- include/ruby/missing.h: fixed merge mistake of r39985. [Backport #8080]

History

#1 - 03/13/2013 11:23 AM - normalperson (Eric Wong)

"jonleighton (Jon Leighton)" j@jonathanleighton.com wrote:

> I have experienced a segfault with Ruby 2 during an IO.select call:
>
> See https://travis-ci.org/jonleighton/spring/jobs/5393025 or https://gist.github.com/jonleighton/5147785 to see the crash output.
>
> I cannot reproduce on a different version of Linux (Fedora). However I was able to reproduce by downloading a VM image of the Travis CI environment and running the code on there (see http://pivotallabs.com/debugging-travis-builds/ for how to do that).
>
> I tried to produce a simple script to reproduce, but without success. I also tried to build Ruby 2 with debugging symbols, but this did not produce the crash. I'm not sure why - perhaps related to compiler optimisations.
>
> I found a workaround for the crash with https://github.com/jonleighton/spring/commit/c8a7afdd3238ef88bffc2c8f56baa21042400e15.

Looking at your workaround, I think a better one is "watcher.to_io"
method needs to memoize its return value.

I think Ruby expects the return value of obj.to_io to be persistent
for the lifetime of obj.

Making IO.select cache the value of to_io internally might be alright...

#2 - 03/13/2013 11:39 AM - kosaki (Motohiro KOSAKI)

- Category set to core

- Status changed from Open to Assigned

- Assignee set to kosaki (Motohiro KOSAKI)

- Target version set to 2.1.0

#3 - 03/16/2013 04:09 AM - jonleighton (Jon Leighton)

normalperson (Eric Wong) wrote:

> Looking at your workaround, I think a better one is "watcher.to_io"
> method needs to memoize its return value.
>
> I think Ruby expects the return value of obj.to_io to be persistent
> for the lifetime of obj.
>
> Making IO.select cache the value of to_io internally might be alright...

It seems that IO.select *does* cache the value of to_io internally. At least that seems to be the case from https://gist.github.com/jonleighton/5172263.

Running the script prints "to_io" once and then hangs.

#4 - 03/16/2013 08:23 AM - normalperson (Eric Wong)

"jonleighton (Jon Leighton)" j@jonathanleighton.com wrote:

> Issue #8080 has been updated by jonleighton (Jon Leighton).
> normalperson (Eric Wong) wrote:
>
> > Looking at your workaround, I think a better one is "watcher.to_io"
> > method needs to memoize its return value.
> >
> > I think Ruby expects the return value of obj.to_io to be persistent
> > for the lifetime of obj.
> >
> > Making IO.select cache the value of to_io internally might be alright...
>
> It seems that IO.select *does* cache the value of to_io internally. At least that seems to be the case from
> https://gist.github.com/jonleighton/5172263.
>
> Running the script prints "to_io" once and then hangs.

That's because nothing else is running while IO.select is running.
IO.select does not cache, it accesses once the data once.
See comments below:

```
----------8<----------
class Foo
def to_io
puts "to_io"
IO.pipe.first
end
end
n = 0

trap(:USR1) { n += 1 }

# Generate garbage to trigger GC, and then EINTR for select()
Thread.new do
loop do
# make some garbage here
(1..1000000).each { |z| z.to_s.dup }
puts "KILL #{n}"
Process.kill("USR1", $$)
end
end

# This will now return empty arrays
# If you swap Foo.new for $stdin, this will never return (as expected)
p IO.select([Foo.new])
----------8<----------
```

#5 - 03/16/2013 02:07 PM - kosaki (Motohiro KOSAKI)

- Status changed from Assigned to Closed

- % Done changed from 0 to 100

This issue was solved with changeset r39775.
Jon, thank you for reporting this issue.
Your contribution to Ruby is greatly appreciated.
May Ruby be with you.

- thread.c: disabled _FORTIFY_SOURCE for avoid to hit glibc bug. [Bug #8080] [ruby-core:53349]
- test/ruby/test_io.rb (TestIO#test_io_select_with_many_files): test for the above.

#6 - 03/16/2013 02:14 PM - kosaki (Motohiro KOSAKI)

- Tracker changed from Bug to Backport

- Project changed from Ruby master to Backport200

- Category deleted (core)

- Status changed from Closed to Assigned

- Assignee changed from kosaki (Motohiro KOSAKI) to nagachika (Tomoyuki Chikanaga)

- Priority changed from Normal to 5

- Target version deleted (2.1.0)

The root cause is, Linux's select implementation supports >1024 files, but glibc doesn't. glibc doesn't correctly understand linux select(2) spec.
This issue is only happen when _FORTIFY_SOURCE >= 1. That said, 2.0 and 1.9.x on Ubuntu. (because Ubuntu enable FORTIFY_SOURCE by default).

I think this fix is important for server usecase. Please backport r39772-r39775.

#7 - 03/16/2013 02:15 PM - kosaki (Motohiro KOSAKI)

The root cause is, Linux's select implementation supports >1024 files, but glibc doesn't. glibc doesn't correctly understand linux select(2) spec.
This issue is only happen when _FORTIFY_SOURCE >= 1. That said, 2.0 and 1.9.x on Ubuntu. (because Ubuntu enable FORTIFY_SOURCE by default).

I think this fix is important for server usecase. Please backport r39772-r39775.

#8 - 03/17/2013 01:07 AM - kosaki (Motohiro KOSAKI)

Hi nagachika-sanm

You need backport 39777, 39779, 39781 and 39783 too.

#9 - 03/20/2013 10:34 PM - nagachika (Tomoyuki Chikanaga)

- Status changed from Assigned to Closed

This issue was solved with changeset r39838.
Jon, thank you for reporting this issue.
Your contribution to Ruby is greatly appreciated.
May Ruby be with you.

---

merge revision(s) 39772,39773: [Backport #8080]

```
* configure.in: check struct timeval exist or not.

* include/ruby/missing.h (struct timeval): check HAVE_STRUCT_TIMEVAL
  properly. and don't include sys/time.h if struct timeval exist.

* file.c: include sys/time.h explicitly.

* random.c: ditto.

* thread_pthread.c: ditto.

* time.c: ditto.

* ext/date/date_strftime.c: ditto.

* include/ruby/missing.h (struct timespec): include <sys/time.h>
```

#10 - 03/20/2013 11:10 PM - nagachika (Tomoyuki Chikanaga)

r39772-r39775, r39777, r39779, r39781 and r39783 are backported.
I also merged r39160, r39162, r39174, r39198 and r39200 for clean merge of configure.in.

#11 - 03/24/2013 05:33 AM - kosaki (Motohiro KOSAKI)

- Project changed from Backport200 to Backport193

- Status changed from Closed to Assigned

- Assignee changed from nagachika (Tomoyuki Chikanaga) to usa (Usaku NAKAMURA)

#12 - 03/28/2013 07:10 PM - usa (Usaku NAKAMURA)

- Status changed from Assigned to Closed

This issue was solved with changeset r39985.
Jon, thank you for reporting this issue.
Your contribution to Ruby is greatly appreciated.
May Ruby be with you.

---

merge revision(s) 39772,39773,39774,39775,39777,39779,39781,39783: [Backport #8080]

```
* configure.in: check struct timeval exist or not.

* include/ruby/missing.h (struct timeval): check HAVE_STRUCT_TIMEVAL
  properly. and don't include sys/time.h if struct timeval exist.

* file.c: include sys/time.h explicitly.

* random.c: ditto.

* thread_pthread.c: ditto.

* time.c: ditto.

* ext/date/date_strftime.c: ditto.

* include/ruby/missing.h (struct timespec): include <sys/time.h>

* include/ruby/missing.h (__syscall): moved to...

* io.c: here. because __syscall() is only used from io.c.
```

```
* include/ruby/missing.h: move "#include <sys/type.h>" to ....

* include/ruby/intern.h: here. because it was introduced for
  fixing NFDBITS issue. [ruby-core:05179].

* thread.c: disabled _FORTIFY_SOURCE for avoid to hit glibc bug.
  [Bug #8080] [ruby-core:53349]

* test/ruby/test_io.rb (TestIO#test_io_select_with_many_files):
  test for the above.

* include/ruby/missing.h: removed __linux__. it's unnecessary.
```

#13 - 03/28/2013 11:08 PM - znz (Kazuhiro NISHIYAMA)

- Status changed from Closed to Assigned

This backport to ruby_1_9_3 is not enough.
Build failed on Mac.
see https://gist.github.com/hsbt/5263190

#14 - 03/29/2013 01:23 PM - usa (Usaku NAKAMURA)

- Status changed from Assigned to Closed

This issue was solved with changeset r39995.
Jon, thank you for reporting this issue.
Your contribution to Ruby is greatly appreciated.
May Ruby be with you.

- include/ruby/missing.h: fixed merge mistake of r39985. [Backport #8080]